

A Fast Summation Method for translation invariant kernels

Fabien Casenave*

August 25, 2015

Abstract

We derive a Fast Multipole Method (FMM) where a low-rank approximation of the kernel is obtained using the Empirical Interpolation Method (EIM). Contrary to classical interpolation-based FMM, where the interpolation points and basis are fixed beforehand, the EIM is a nonlinear approximation method which constructs interpolation points and basis which are adapted to the kernel under consideration. The basis functions are obtained using evaluations of the kernel itself. We restrict ourselves to translation-invariant kernels, for which a modified version of the EIM approximation can be used in a multilevel FMM context; we call the obtained algorithm Empirical Interpolation Fast Multipole Method (EIFMM). An important feature of the EIFMM is a built-in error estimation of the interpolation error made by the low-rank approximation of the far-field behavior of the kernel: the algorithm selects the optimal number of interpolation points required to ensure a given accuracy for the result, leading to important gains for inhomogeneous kernels.

1 Introduction

We consider the following problem: compute approximations of many sums of the form

$$f(\bar{x}_i) = \sum_{j=1}^N \sigma_j K(\bar{x}_i, \bar{y}_j), 1 \leq i \leq N, \quad (1)$$

for a set of potentials σ_j , and target points $\bar{x}_i \in \Omega$ and source points $\bar{y}_j \in \Omega$, where $\Omega \subset \mathbb{R}^3$, with a controllable error ϵ . The domain Ω and N are fixed for all computations, but the potentials σ_j and the points \bar{x}_i and \bar{y}_j can vary from one computation to another. The many-query situation can come from a temporal N-body problem (in which case the set of source points equals the set of observation points), or the iterative resolution of a dense linear system.

Various algorithms have been proposed to tackle such problems. Some methods only using numerical evaluations of K have been proposed, see [1, 8, 15]. Many methods, e.g. the mosaic skeleton method [27], hierarchical matrices [20], the Fast Multipole Method (FMM) [17] (and its algebraic generalization H2-matrices [18, 19]), the panel clustering method [21], rely on local low-rank approximations of K :

$$K(x, y) \approx \sum_{l,m=1}^d \Delta_{l,m} g_l(x) h_m(y), x \in I, y \in J, \quad (2)$$

where I and J are subsets of Ω . Such approximations can be analytic, like interpolation (for instance the Chebyshev interpolation as in [14, 20]), or multipole expansions. FMM for general kernels have been developed leading to kernel-independent procedures, see [24, 28]. In [13], another kernel-independent FMM algorithm, based on the Chebyshev interpolation to obtain the low-rank approximation (2) of K , is developed and called Black-Box FMM. Since the algorithm we present here is in the same family as the one developed in [13], our numerical experiments contain comparisons with this method. For other kernel-independent interpolation based procedures, see [10, 11]. Adaptive Cross Approximation (ACA) [3, 4] directly constructs approximations from the evaluations of the kernel on the target and source points $K(\bar{x}_i, \bar{y}_j)$, by selecting some columns and rows of the matrix $(K(\bar{x}_i, \bar{y}_j))_{i,j}$.

*IGN LAREG, Univ Paris Diderot, Sorbonne Paris Cité, 5 rue Thomas Mann, 75205 Paris Cedex 13, France.

Approximating a matrix by extracting some of its rows and columns has been investigated in [16]. With ACA methods, (2) is such that g_l depends on J and h_m depends on I , which leads to a storage complexity of $dN \log(N)$. With FMM and H2-matrices, g_l and h_m are independent of respectively J and I , which allows additional factorization and the use of nested basis to lower the storage complexity to dN . An improved ACA exploiting the nested basis property has been proposed in [5].

In this work, a low-rank approximation of K is obtained with the Empirical Interpolation Method (EIM), a procedure that represents a two-variable function in a so-called affine dependent formula and that is used in the Reduced Basis community, [2]. The EIM provides a non-linear and function-dependent interpolation: the basis functions and interpolation points are not fixed a priori but are computed during a greedy procedure, which differs from the Chebyshev interpolations. The procedure should capture any irregularity of the kernel better than a fixed basis interpolation procedure, since the basis functions of the EIM are based on evaluations of the kernel itself. Theoretical results on the convergence of EIM interpolation error in the general case are limited, to the author knowledge, to an upper bound of $2^d - 1$ of the Lebesgue constant, where d is the number of interpolation points, see [2, Proposition 3.2]. Various numerical studies illustrate the efficiency of the EIM [12, 22, 23, 26], even though a theoretical analysis validating this efficiency is not available yet. Depending on the kernel (and on the level in the tree for inhomogeneous kernels), the EIM selects an optimal number of terms that ensures a given accuracy of the approximation. Such a variable-order interpolation procedure have been used in [6] in the H2-matrix context with Lagrange polynomials.

A direct application of the EIM algorithm to the kernel K to obtain a local low-rank approximation (2) is such that g_l depends on J and h_m depends on I , and therefore is not compatible with a multilevel FMM: a nested basis procedure cannot be derived. To correct this, we make use of the translation-invariant hypothesis of K and derive a local-global low-rank EIM approximation of K , in the sense that one variable is restricted to a local small subdomain while the other variable is taken in almost all the domain. Then, we derive recursive formulae to produce a multilevel FMM for this modified approximation: this is detailed in Section 3.3. The modified approximation uses four summations, instead of two summations for the Black-Box FMM. However, we will see that the M2L step, the most expensive one, has the same complexity as the Black-Box FMM one with respect to the number on interpolation points. Our algorithm requires a precomputation step, during which the greedy step of the EIM at each level is carried-out, at well as compression of some operators. In our simulations, we did not notice any speedup on homogeneous kernels compared to Black-Box FMM. For inhomogeneous kernels, the approximation problem gets easier as we get higher in the tree: corresponding EIM require much less interpolation points and speedups are measured. In 10^6 points test-cases, depending on the kernel, the EIFMM gets more interesting if one has to compute more than 2 to 18 queries of the summation. In a 10^8 points test-case, for the kernel e^{-r^2} , the precomputation step becomes negligible, and a single query is computed twice as fast for a nearly one order of magnitude better accuracy.

In Section 2, the EIM is recalled. In Section 3 is presented the new FMM algorithm, called the Empirical Interpolation Fast Multipole Method (EIFMM): first, a four-summation formula based on the EIM is derived in Section 3.1, that can be used in a multilevel FMM procedure, then a monolevel version of the EIFMM is proposed in Section 3.2, and a multilevel one in Section 3.3. In Section 4, the overall execution time is reduced using some precomputation and classical compression of the M2L step. Finally, numerical experiments are presented in Section 5, and conclusions are drawn in Section 6.

2 The Empirical Interpolation Method

Consider a function $K(x, y)$ defined over $\mathcal{D}_x \times \mathcal{D}_y$ assumed to be real-valued for simplicity, where $\mathcal{D}_x, \mathcal{D}_y \subset \Omega$. Fix an integer d . The Empirical Interpolation Method provides a way to approximate this function in the following form:

$$K(x, y) \approx (I_d K)(x, y) := \sum_{m=1}^d \lambda_m(x) q_m(y), \quad (3)$$

where $\lambda_m(x)$ is such that

$$\sum_{m=1}^d B_{l,m} \lambda_m(x) = K(x, y_l), \quad \forall 1 \leq l \leq d. \quad (4)$$

The functions $q_m(\cdot)$ and the matrix $B \in \mathbb{R}^{d \times d}$, which is lower triangular with unity diagonal, are constructed in Algorithm 1, where $\delta_d = \text{Id} - I_d$ and $\|\cdot\|_{\mathcal{D}_y}$ is a norm on \mathcal{D}_y , for instance the $L^\infty(\mathcal{D}_y)$ - or the $L^2(\mathcal{D}_y)$ -norm. In practice, the argmax appearing in Algorithm 1 is searched over finite subsets of \mathcal{D}_x and \mathcal{D}_y , denoted respectively by $\mathcal{D}_{x,\text{trial}}$ and $\mathcal{D}_{y,\text{trial}}$ and called training sets. Notice that Algorithm 1 also constructs the set of points $\{y_l\}_{1 \leq l \leq d}$ in \mathcal{D}_y used in (4), and a set of points $\{x_l\}_{1 \leq l \leq d}$ in \mathcal{D}_x . These points are interpolation points for the approximation of K (see Property 2.1) and are different from the target and source points \bar{x}_i, \bar{y}_j .

Algorithm 1 Greedy algorithm of the EIM

1. Choose $d > 1$
 2. Set $k := 1$
 3. Compute $x_1 := \underset{x \in \mathcal{D}_x}{\operatorname{argmax}} \|K(x, \cdot)\|_{\mathcal{D}_y}$
 4. Compute $y_1 := \underset{y \in \mathcal{D}_y}{\operatorname{argmax}} |K(x_1, y)|$ [First interpolation point]
 5. Set $q_1(\cdot) := \frac{K(x_1, \cdot)}{K(x_1, y_1)}$ [First basis function]
 6. Set $B_{1,1} := 1$ [Initialize matrix B]
 7. **while** $k < d$ **do**
 8. Compute $x_{k+1} := \underset{x \in \mathcal{D}_x}{\operatorname{argmax}} \|(\delta_k K)(x, \cdot)\|_{\mathcal{D}_y}$
 9. Compute $y_{k+1} := \underset{y \in \mathcal{D}_y}{\operatorname{argmax}} |(\delta_k K)(x_{k+1}, y)|$ [($k+1$)-th interpolation point]
 10. Set $q_{k+1}(\cdot) := \frac{(\delta_k K)(x_{k+1}, \cdot)}{(\delta_k K)(x_{k+1}, y_{k+1})}$ [($k+1$)-th basis function]
 11. Set $B_{k+1,m} := q_m(y_{k+1})$, for all $1 \leq m \leq k+1$ [Increment matrix B]
 12. $k \leftarrow k+1$ [Increment the size of the decomposition]
 13. **end while**
-

The following assumption is made:

(H) The dimension of $\text{Span}(K(x, \cdot), \cdot)_{x \in \mathcal{D}_x}$ is larger than d .

From [(H)], the functions $\{K(x_l, \cdot)\}_{1 \leq l \leq d}$ are linearly independent (otherwise, $(\delta_k K)(\mu_{k+1}, x_{k+1}) = 0$ for some k in Algorithm 1).

It is shown in [7] that the approximation (3) can be written in the form

$$(I_d K)(x, y) = \sum_{l,m=1}^d \Delta_{l,m} K(x, y_l) K(x_m, y), \quad (5)$$

where $\Delta = B^{-t} \Gamma^{-1}$, and where the matrix Γ is upper triangular and constructed recursively in the loop in k of Algorithm 1 in the following way:

- $k = 1$:

$$\Gamma_{1,1} = K(x_1, y_1),$$

- $k \rightarrow k+1$:

$$\begin{aligned} \Gamma_{k+1,k+1} &= (\delta_k K)(x_{k+1}, y_{k+1}), \\ \Gamma_{m,k+1} &= 0, & \forall 1 \leq m \leq k, \\ \Gamma_{k+1,m} &= \alpha_m, & \forall 1 \leq m \leq k, \end{aligned}$$

where the vector α is such that $\sum_{m=1}^k B_{l,m} \alpha_m = K(x_{k+1}, y_l)$, for all $1 \leq l \leq k$.

The particular form (5) is the key to the use of the EIM approximation in a fast multipole context.

We recall the interpolation property of $I_d K$; see [23, Lemma 1]:

Property 2.1 (Interpolation property). *For all $1 \leq m \leq d$,*

$$\begin{cases} (I_d K)(x, y_m) = K(x, y_m), & \text{for all } x \in \mathcal{D}_x, \\ (I_d K)(x_m, y) = K(x_m, y), & \text{for all } y \in \mathcal{D}_y. \end{cases}$$

An important aspect is that the error made by the EIM approximation over the training set $\mathcal{D}_{x,\text{trial}} \times \mathcal{D}_{y,\text{trial}}$, denoted ϵ_k and defined as $\epsilon_k := \max_{y \in \mathcal{D}_y} |(\delta_k K)(x_{k+1}, y)|$, is monitored during the greedy procedure. In practice, a stopping criterion is used on ϵ_k instead of on the number of interpolation points d , which was used to simplify the presentation of the EIM. Hence, depending on the regularity of the function K , the EIM will automatically select the number of interpolation points to certify the error over $\mathcal{D}_{x,\text{trial}} \times \mathcal{D}_{y,\text{trial}}$.

3 The Empirical Interpolation Fast Multipole Method

The FMM algorithm is based on a low-rank approximation of the far-field behavior of the kernel of the form (2). Such approximations enable some factorizations in the computation of key-quantities and the propagation of information through a tree structure, resulting in lowering the complexity of the computation of (1). For a general presentation of the FMM algorithm, see the original article [17, Section 3].

3.1 A suitable approximation

For the simplicity of the presentation, the algorithm is described in 1D. It can be readily extended in 2D and 3D. Consider a kernel $K(x, y)$, $(x, y) \in \Omega \times \Omega$, where $\Omega = (-\frac{L}{2}, \frac{L}{2})$ is the complete domain on which the FMM computation is carried out. The kernel is supposed translation invariant, i.e. such that $K(x, y) = K(x - a, y - a)$, $\forall a \in \mathbb{R}$. The domain Ω is partitioned in 2^κ intervals, $\kappa \in \mathbb{N}^*$, of length $2l_\kappa$, with $l_\kappa := 2^{-\kappa-1}L$. Consider two intervals I and J , well-separated in the sense that $\min_{(x,y) \in I \times J} |x - y| \geq 2l_\kappa$. To anticipate the generalization to higher dimensions, I and J are called “boxes” in what follows. We denote c_I and c_J the center of the boxes I and J . Consider an EIM approximation constructed over $I \times J$:

$$K(x, y) \approx \sum_{l,m=1}^d \Delta_{l,m}^{I,J} K(x, y_l^J) K(x_m^I, y), \quad x \in I, \quad y \in J, \quad (6)$$

where I and J superscripts have been added to underline the box-dependent quantities (as deduced from Algorithm 1). In (6), $K(x, y_l^J)$ depends on J , the box containing y and $K(x_m^I, y)$ depends on I , the box containing x : as stated in the introduction, this approximation is then not suitable for a multilevel FMM procedure.

Define now $I_0 := \{x - c_{\hat{J}}, x \in \hat{I}\}$ for all box \hat{J} in the partitioning of Ω and all box \hat{I} well separated from \hat{J} , and define J_0 the box of length $2l_\kappa$ located at the center of Ω (notice that J_0 is not a box from the partitioning of Ω). In our setting, $J_0 = (-l_\kappa, l_\kappa)$ and $I_0 = (-L + l_\kappa, -3l_\kappa) \cup (3l_\kappa, L - l_\kappa)$, see Figure 1. In other words, J_0 is a source box translated to the center of the domain, I_0 is the union of possible target boxes, translated in the same fashion.

Since $K(x, y) = K(x - c_J, y - c_J)$, with $x - c_J \in I_0$ and $y - c_J \in J_0$, it is inferred that

$$K(x, y) \approx \sum_{l,m=1}^d \Delta_{l,m}^{I_0, J_0} K(x - c_J, y_l^{J_0}) K(x_m^{I_0}, y - c_J), \quad (7)$$

where the EIM is carried out on $\mathcal{D}_x = I_0$ and $\mathcal{D}_y = J_0$. An important aspect is that I_0 and J_0 are fixed domains, independent from the boxes I and J containing respectively x and y . Consider the




Figure 1: Representation of the partitioning of Ω in 1D for $\kappa = 3$ and of the corresponding boxes I_0 and J_0 .

first evaluation of the kernel in (7). There holds $K(x - c_J, y_l^{J_0}) = K(x - c_I, y_l^{J_0} + c_J - c_I)$, where $x - c_I \in J_0$ and $y_l^{J_0} + c_J - c_I \in I_0$. It is then inferred that

$$K(x - c_J, y_l^{J_0}) \approx \sum_{l', m'=1}^d \Delta_{l', m'}^{J_0, I_0} K(x - c_I, y_{l'}^{I_0}) K(x_{m'}^{J_0}, y_l^{J_0} + c_J - c_I), \quad 1 \leq l \leq m, \quad (8)$$

where a second EIM has been carried out on $\mathcal{D}_x = J_0$ and $\mathcal{D}_y = I_0$ (with the same number of terms d for simplicity). Notice that if the kernel is symmetric, $\Delta_{l', m'}^{I_0, J_0} = (\Delta_{l', m'}^{J_0, I_0})^t$, $x_m^{I_0} = y_m^{I_0}$ for all $1 \leq m \leq d$, and $y_m^{J_0} = x_m^{J_0}$ for all $1 \leq m \leq d$. Injecting (8) into (7):

$$K(x, y) \approx \sum_{l'=1}^d K(x - c_I, y_{l'}^{I_0}) \sum_{m'=1}^d \Delta_{l', m'}^{J_0, I_0} \sum_{l=1}^d K(x_{m'}^{J_0}, y_l^{J_0} + c_J - c_I) \sum_{m=1}^d \Delta_{l, m}^{I_0, J_0} K(x_m^{I_0}, y - c_J), \quad x \in I, y \in J. \quad (9)$$

In (9), the only term depending on the boxes I and J (besides x and y themselves) are c_I and c_J , the center of the boxes I and J . Hence, in the low-rank approximation, the x -dependent function $K(x - c_I, y_{l'}^{I_0})$ no longer depends on J and the y -dependent function $K(x_m^{I_0}, y - c_J)$ no longer depends on I . As a consequence, the non-trivial manipulation leading to (9) are required for the use of EIM interpolation in a multilevel FMM context.

Notice that in all the kernel evaluations in (9), the first and second variables are always separated by a distance of at least $2l_\kappa$.

3.2 A monolevel multipole method

Consider a set of N target points $\{\bar{x}_i\}_{1 \leq i \leq N}$, N source points $\{\bar{y}_j\}_{1 \leq j \leq N}$ in Ω , and N potentials $\{\sigma_j\}_{1 \leq j \leq N}$, with $N \gg 1$. We look for a low complexity approximation of

$$f(\bar{x}_i) = \sum_{j=1}^N \sigma_j K(\bar{x}_i, \bar{y}_j), \quad 1 \leq i \leq N. \quad (10)$$

We say that two boxes I and J are well-separated if $\min_{x, y \in I \times J} |x - y| \geq 2l_\kappa$. The set of boxes well-separated from I is denoted $\mathcal{F}(I)$, its complement in Ω is the neighborhood of I and is denoted $\mathcal{N}(I)$. The FMM aims at compressing the “far” interactions in the following fashion:

$$f_{\mathcal{F}}(\bar{x}_i) = \sum_{J \in \mathcal{F}(I)} \sum_{\bar{y}_j \in J} \sigma_j K(\bar{x}_i, \bar{y}_j), \quad \text{for all box } I \text{ and all } \bar{x}_i \in I. \quad (11)$$

A monolevel multipole method can be directly derived from (9) to approximate $f_{\mathcal{F}}(\bar{x}_i)$, see Algorithm 2. The obtained approximation of $f_{\mathcal{F}}(\bar{x}_i)$ is denoted $f_{\mathcal{F}, \text{FMM}}(\bar{x}_i)$. An approximation of $f(\bar{x}_i)$ is then given by

$$f(\bar{x}_i) \approx f_{\mathcal{F}, \text{FMM}}(\bar{x}_i) + f_{\mathcal{N}}(\bar{x}_i), \quad (12)$$

Step number	Complexity
1	dN
2	$d^2(2^D)^\kappa$
3	$d^2(2^D)^{2\kappa}$
4	$d^2(2^D)^\kappa$
5	dN
f_N	$N^2(2^D)^{-\kappa}$

Table 1: Complexity of each step of Algorithm 2 and of the near interactions term.

where the “near” interactions are defined by

$$f_N(\bar{x}_i) = \sum_{\bar{y}_j \in \mathcal{N}(I)} \sigma_j K(\bar{x}_i, \bar{y}_j). \quad (13)$$

Algorithm 2 Monolevel EIFMM

1. Compute $W_m^I := \sum_{\bar{y}_j \in I} \sigma_j K(x_m^{I_0}, \bar{y}_j - c_I)$, for all box I and all $1 \leq m \leq d$
 2. Compute $\hat{W}_l^I := \sum_{m=1}^d \Delta_{l,m}^{I_0, J_0} W_m^I$, for all box I and all $1 \leq l \leq d$
 3. Compute $l_{m'}^I := \sum_{J \in \mathcal{F}(I)} \sum_{l=1}^d K(x_{m'}^{J_0}, y_l^{J_0} + c_J - c_I) \hat{W}_l^J$, for all box I and all $1 \leq m' \leq d$
 4. Compute $\hat{l}_{l'}^I := \sum_{m'=1}^d \Delta_{l',m'}^{J_0, I_0} l_{m'}^I$, for all box I and all $1 \leq m' \leq d$
 5. Compute $f_{\mathcal{F}, \text{FMM}}(\bar{x}_i) := \sum_{l'=1}^d K(\bar{x}_i - c_I, y_{l'}^{I_0}) \hat{l}_{l'}^I$, for all box I and all $x_i \in I$
-

In Table 1 is given the complexity of each step of Algorithm 2. For step 3 and the computation of the near interactions to be in the same order in N , one must impose $(2^D)^{2\kappa} \propto N^2(2^D)^{-\kappa}$, leading to a overall complexity of the sum of order $N^{\frac{4}{3}}$. As classically done in the literature, the derivation of a multilevel multipole method allows to lower the complexity in N of the algorithm to linear.

3.3 A multilevel multipole method

The domain of computation $\Omega = (-\frac{L}{2}, \frac{L}{2})$ is now organized as a binary tree with $\kappa + 1$ levels. Level 0 is the root of the tree: there is one box equal to Ω , level 1 contains the partitioning of Ω into two boxes of length $\frac{L}{2}$, and the following levels are obtained by partitioning each box of the previous level in the same fashion, up to level κ .

A superscript k is now added to the quantities previously defined, to indicate that they depend on the level k . Two EIM’s are to be considered at each level, in the same fashion as in Section 3.1, where now the domains depend on the level k in the tree: $I_0^k = (-L + l_k, -3l_k) \cup (3l_k, L - l_k)$, and $J_0^k = (-l_k, l_k)$.

For symmetric kernels, this amount of precomputation can be lowered to one EIM by level. The key to a multilevel method stands in the ability to derive recursion formulae to propagate information from the level κ to the level 0 in the upward pass, and from the level 0 to the level κ in the downward pass.

Black-Box FMM, at least in its original version, is proposed with a constant-order of interpolation, which enables to propagate the information through the levels of the tree without introducing additional errors. Considering a variable-order Chebyshev interpolation is possible, but one must carefully analyse the introduced error, as done in [6] in the context of H2-matrices. In our method, the domains of approximation depicted in Figure 1 naturally enable a variable-order interpolation depending on the level of the tree, for which the introduced error is the same as the EIM approximation error in (9). As a consequence, the variable-order aspect of the approximation is directly obtained at the required accuracy, and the length of the approximation is automatically optimized at each level of the tree, which leads to important gains for inhomogeneous kernels.

3.3.1 Recursion for the upward pass: multipole to multipole (M2M)

The goal is to compute

$$W_m^{I^k} := \sum_{\bar{y}_j \in I^k} \sigma_j K(x_m^{I_0^k}, \bar{y}_j - c_{I^k}), \quad (14)$$

for all box I^k at all level k with $0 \leq k \leq \kappa$ and all $1 \leq m \leq d^k$ by recursion. The initialization of the recursion consists in computing $W_m^{I^\kappa}$ at level κ using (14). Then, suppose $W_m^{I^k}$ is available at level k , there holds

$$W_m^{I^{k-1}} = \sum_{\bar{y}_j \in I^{k-1}} \sigma_j K(x_m^{I_0^{k-1}}, \bar{y}_j - c_{I^{k-1}}) = \sum_{J^k \in \mathcal{C}(I^{k-1})} \sum_{\bar{y}_j \in J^k} \sigma_j K(x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k}, \bar{y}_j - c_{J^k}), \quad (15)$$

where $\mathcal{C}(I^{k-1})$ denotes the children boxes of I^{k-1} . To obtain a recursion on $W_m^{I^{k-1}}$, i.e. to compute $W_m^{I^{k-1}}$ using some values of $W_m^{J^k}$, for J^k being boxes at the level k , we will use the EIM approximations already computed at level k to derive (9). One must check that the corresponding EIM approximations are evaluated on the right domains, namely I_0^k and J_0^k .

For all $\bar{y}_j \in J^k$, $-l_k \leq \bar{y}_j - c_{J^k} \leq l_k$, i.e. $\bar{y}_j - c_{J^k} \in J_0^k$. Then, for all $J^k \in \mathcal{C}(I^{k-1})$, $-l_k \leq c_{I^{k-1}} - c_{J^k} \leq l_k$ (see for instance Figure 3). Moreover, since $x_m^{I_0^{k-1}} \in I_0^{k-1} = (-L + l_{k-1}, -3l_{k-1}) \cup (3l_{k-1}, L - l_{k-1})$, it is inferred that

$$-L + l_k \leq x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k} \leq -5l_k \leq -3l_k \text{ or } 3l_k \leq 5l_k \leq x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k} \leq L - l_k,$$

i.e. $x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k} \in I_0^k$. Therefore, the EIM from level k can then be used to write the following approximation:

$$K(x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k}, \bar{y}_j - c_{J^k}) \approx \sum_{p,q=1}^{d^k} \Delta_{p,q}^{I_0^k, J_0^k} K(x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k}, y_p^{J_0^k}) K(x_q^{I_0^k}, \bar{y}_j - c_{J^k}), \quad (16)$$

where the approximation error is of the same order as in the (9).

Then, injecting (16) into (15), it is inferred that

$$\begin{aligned} W_m^{I^{k-1}} &\approx \sum_{p,q=1}^{d^k} \Delta_{p,q}^{I_0^k, J_0^k} \sum_{J^k \in \mathcal{C}(I^{k-1})} K(x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k}, y_p^{J_0^k}) \sum_{\bar{y}_j \in J^k} \sigma_j K(x_q^{I_0^k}, \bar{y}_j - c_{J^k}) \\ &\approx \sum_{p,q=1}^{d^k} \Delta_{p,q}^{I_0^k, J_0^k} \sum_{J^k \in \mathcal{C}(I^{k-1})} K(x_m^{I_0^{k-1}} + c_{I^{k-1}} - c_{J^k}, y_p^{J_0^k}) W_q^{J^k}, \end{aligned} \quad (17)$$

which provides a recursive formula for the upward pass. In the same fashion as for the monolevel version, we then compute

$$\hat{W}_l^{I^k} := \sum_{m=1}^{d^k} \Delta_{l,m}^{I_0^k, J_0^k} W_m^{I^k}, \quad (18)$$

for all box I^k at all level k with $0 \leq k \leq \kappa$ and all $1 \leq l \leq d^k$.

3.3.2 Transfer pass: multipole to local (M2L)

We denote $\mathcal{I}(I^k)$, the set of boxes that are the children of the boxes in the neighborhood of the parent of I^k but are well-separated from I^k at level k , and call it the interaction list of I^k , see Figure 2.

The M2L step consists in computing

$$g_{m'}^{I^k} := \sum_{J^k \in \mathcal{I}(I^k)} \sum_{l=1}^{d^k} K(x_{m'}^{J_0^k}, y_l^{J_0^k} + c_{J^k} - c_{I^k}) \hat{W}_l^{J^k}, \quad (19)$$

for all box I^k at all level k with $0 \leq k \leq \kappa$ and all $1 \leq m' \leq d^k$.




Figure 2: Representation of the neighborhood and the interaction list of a box I^k and level k .

3.3.3 Recursion for the downward pass: local to local (L2L)

The goal is to compute

$$l_{m'}^{I^k} := \sum_{J^k \in \mathcal{F}(I^k)} \sum_{l=1}^{d^k} K(x_{m'}^{J_0^k}, y_l^{J_0^k} + c_{J^k} - c_{I^k}) \hat{W}_l^{J^k}, \quad (20)$$

for all box I^k at all level k with $0 \leq k \leq \kappa$ and all $1 \leq m' \leq d^k$ by recursion. The initialization of the recursion consists in computing $l_{m'}^{I^2}$ at level 2 using (20). Then, suppose $l_{m'}^{I^k}$ is available, there holds

$$\begin{aligned} l_{m'}^{I^{k+1}} &= g_{m'}^{I^{k+1}} + \sum_{J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))} \sum_{J^{k+1} \in \mathcal{C}(J^k)} \sum_{l=1}^{d^{k+1}} K(x_{m'}^{J_0^{k+1}}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{I^{k+1}}) \hat{W}_l^{J^{k+1}} \\ &= g_{m'}^{I^{k+1}} + \sum_{J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))} \sum_{J^{k+1} \in \mathcal{C}(J^k)} \sum_{l=1}^{d^{k+1}} K(x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}) \hat{W}_l^{J^{k+1}}, \end{aligned} \quad (21)$$

where $\mathcal{P}(I^{k+1})$ denotes the parent of I^{k+1} .

It is seen on Figures 3 and 4 that

$$-l_{k+1} \leq c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})} \leq l_{k+1},$$

and that

$$-L + 3l_{k+1} \leq c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})} \leq -7l_{k+1} \text{ or } 7l_{k+1} \leq c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})} \leq L - 3l_{k+1}.$$




Figure 3: Possible values for $c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}$ and $c_{J^{k+1}} - c_{J^k}$, and smallest possible values for $c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}$ and $c_{\mathcal{P}(I^{k+1})} - c_{J^k}$, with J^k well-separated from $\mathcal{P}(I^{k+1})$ at level k , and J^{k+1} child of J^k .

Since $x_{m'}^{J_0^{k+1}}, y_l^{J_0^{k+1}} \in J_0^{k+1}$, it is then readily verified that $x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})} \in J_0^k$ and

$$-L + 2l_{k+1} = -L + l_k \leq y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})} \leq -6l_{k+1} = -3l_k \text{ or } 3l_k \leq y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})} \leq L - l_k,$$




Figure 4: Largest possible value for $c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}$ and $c_{\mathcal{P}(I^{k+1})} - c_{J^k}$, with J^k well-separated from $\mathcal{P}(I^{k+1})$ at level k , and J^{k+1} child of J^k .

i.e. $y_l^{J^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})} \in I_0^k$. Therefore, the EIM from level k can then be used to write the following approximation

$$\begin{aligned} & K(x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}) \approx \\ & \sum_{p', q'=1}^{d^k} \Delta_{p', q'}^{J_0^k, I_0^k} K(x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}, y_{p'}^{I_0^k}) K(x_{q'}^{J_0^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}). \end{aligned} \quad (22)$$

Then, injecting (22) into (21), it is inferred that

$$\begin{aligned} l_{m'}^{I^{k+1}} & \approx g_{m'}^{I^{k+1}} + \sum_{J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))} \sum_{J^{k+1} \in \mathcal{C}(J^k)} \sum_{p', q'=1}^{d^k} \Delta_{p', q'}^{J_0^k, I_0^k} \sum_{l=1}^{d^{k+1}} K(x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}, y_{p'}^{I_0^k}) \\ & K(x_{q'}^{J_0^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}) \hat{W}_l^{J^{k+1}}. \end{aligned} \quad (23)$$

Consider $K(x_{q'}^{J_0^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{\mathcal{P}(I^{k+1})}) = K(x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{J^k})$. Since $J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))$ and J^k is the parent of J^{k+1} , it is seen on Figures 3 and 4 that

$$-l_{k+1} \leq c_{J^{k+1}} - c_{J^k} \leq l_{k+1},$$

and that

$$-L + 4l_{k+1} \leq c_{\mathcal{P}(I^{k+1})} - c_{J^k} \leq -8l_{k+1} \text{ or } 8l_{k+1} \leq c_{\mathcal{P}(I^{k+1})} - c_{J^k} \leq L - 4l_{k+1}.$$

Since $x_{q'}^{J_0^k}, y_l^{J_0^{k+1}} \in J_0^{k+1}$, $J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))$, and J^k is the parent of J^{k+1} , it is readily verified that $y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{J^k} \in J_0^k$ and

$$-L + l_k \leq -L + 3l_{k+1} \leq x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k} \leq -7l_{k+1} \leq -3l_k \text{ or } 3l_k \leq x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k} \leq L - l_k,$$

i.e. $x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k} \in I_0^k$. Therefore, the EIM from level k can then be used again to write the following approximation

$$\begin{aligned} & K(x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{J^k}) \approx \\ & \sum_{u, v=1}^{d^k} \Delta_{u, v}^{I_0^k, J_0^k} K(x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k}, y_u^{J_0^k}) K(x_v^{I_0^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{J^k}). \end{aligned} \quad (24)$$

Injection (24) into (23), making use of the formula (18) to express $\hat{W}_l^{J^{k+1}}$, it is inferred that

$$\begin{aligned} l_{m'}^{I^{k+1}} & \approx g_{m'}^{I^{k+1}} + \sum_{J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))} \sum_{J^{k+1} \in \mathcal{C}(J^k)} \sum_{p', q'=1}^{d^k} \Delta_{p', q'}^{J_0^k, I_0^k} \sum_{l=1}^{d^{k+1}} K(x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}, y_{p'}^{I_0^k}) \\ & \sum_{u, v=1}^{d^k} \Delta_{u, v}^{I_0^k, J_0^k} K(x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J^k}, y_u^{J_0^k}) K(x_v^{I_0^k}, y_l^{J_0^{k+1}} + c_{J^{k+1}} - c_{J^k}) \sum_{m=1}^{d^{k+1}} \Delta_{l, m}^{I_0^{k+1}, J_0^{k+1}} W_m^{J^{k+1}}. \end{aligned} \quad (25)$$

Making use of $K(x_v^{I_0^k}, y_l^{J_0^{k+1}} + c_{J_{k+1}} - c_{J_k}) = K(x_v^{I_0^k} + c_{J_k} - c_{J_{k+1}}, y_l^{J_0^{k+1}})$ and $K(x_{q'}^{J_0^k} + c_{\mathcal{P}(I^{k+1})} - c_{J_k}, y_u^{J_0^k}) = K(x_{q'}^{J_0^k}, y_u^{J_0^k} + c_{J_k} - c_{\mathcal{P}(I^{k+1})})$, and reorganizing some terms, there holds

$$\begin{aligned}
l_{m'}^{I^{k+1}} &\approx g_{m'}^{I^{k+1}} + \sum_{p', q'=1}^{d^k} \Delta_{p', q'}^{J_0^k, I_0^k} K(x_{m'}^{J_0^{k+1}} + c_{I^{k+1}} - c_{\mathcal{P}(I^{k+1})}, y_{p'}^{I_0^k}) \\
&\quad \sum_{J^k \in \mathcal{F}(\mathcal{P}(I^{k+1}))} \sum_{u=1}^{d^k} K(x_{q'}^{J_0^k}, y_u^{J_0^k} + c_{J_k} - c_{\mathcal{P}(I^{k+1})}) \underbrace{\sum_{v=1}^{d^k} \Delta_{u,v}^{I_0^k, J_0^k} \sum_{l,m=1}^{d^{k+1}} \Delta_{l,m}^{I_0^{k+1}, J_0^{k+1}} \sum_{J^{k+1} \in \mathcal{C}(J^k)} K(x_v^{I_0^k} + c_{J_k} - c_{J_{k+1}}, y_l^{J_0^{k+1}}) W_m^{J^{k+1}}}_{=W_v^{J^k} \text{ using (17)}} \\
&\quad \underbrace{\hspace{10em}}_{=\hat{W}_u^{J^k} \text{ using (18)}} \\
&\quad \underbrace{\hspace{10em}}_{=l_{q'}^{\mathcal{P}(I^{k+1})} \text{ using (20)}}
\end{aligned} \tag{26}$$

which provides a recursive formula for the upward pass. In the same fashion as for the monolevel version, we then compute

$$\hat{l}_l^{I^\kappa} := \sum_{m'=1}^{d^\kappa} \Delta_{l', m'}^{J_0^\kappa, I_0^\kappa} l_{m'}^{I^\kappa}, \tag{27}$$

for all box I^κ at level κ and all $1 \leq l' \leq d^\kappa$.

3.3.4 Algorithm and complexity

Using the recursion formulae derived above, the multilevel procedure is detailed in Algorithm 3.

Algorithm 3 Multilevel EIFMM

1. Compute $W_m^{I^\kappa} = \sum_{\bar{y}_j \in I^\kappa} \sigma_j K(x_m^{I_0^\kappa}, \bar{y}_j - c_{I^\kappa})$, for all box I^κ at level κ and all $1 \leq m \leq d^\kappa$
 2. Compute $W_m^{I^k} = \sum_{p,q=1}^{d^{k+1}} \Delta_{p,q}^{I_0^{k+1}, J_0^{k+1}} \sum_{J^{k+1} \in \mathcal{C}(I^k)} K(x_m^{I_0^k} + c_{I^k} - c_{J_{k+1}}, y_p^{J_0^{k+1}}) W_q^{J^{k+1}}$, for all box I^k at all levels $\kappa - 1 \geq k \geq 0$ and all $1 \leq m \leq d^{k+1}$
 3. Compute $\hat{W}_l^{I^k} = \sum_{m=1}^{d^k} \Delta_{l,m}^{I_0^k, J_0^k} W_m^{I^k}$, for all box I^k at all level k with $0 \leq k \leq \kappa$ and all $1 \leq l \leq d^k$
 4. Compute $g_{m'}^{I^k} = \sum_{J^k \in \mathcal{I}(I^k)} \sum_{l=1}^{d^k} K(x_{m'}^{J_0^k}, y_l^{J_0^k} + c_{J^k} - c_{I^k}) \hat{W}_l^{J^k}$, for all box I^k at all level k with $0 \leq k \leq \kappa$ and all $1 \leq m' \leq d^k$.
 5. Let $l_{m'}^{I^0} = g_{m'}^{I^0}$ and compute $l_{m'}^{I^k} = g_{m'}^{I^k} + \sum_{p', q'=1}^{d^{k-1}} \Delta_{p', q'}^{J_0^{k-1}, I_0^{k-1}} K(x_{m'}^{J_0^k} + c_{I^k} - c_{\mathcal{P}(I^k)}, y_{p'}^{I_0^{k-1}}) l_{q'}^{\mathcal{P}(I^k)}$ for all box I^k at all levels $1 \leq k \leq \kappa$ and all $1 \leq m' \leq d^{k-1}$
 6. Compute $\hat{l}_{l'}^{I^\kappa} := \sum_{m'=1}^{d^\kappa} \Delta_{l', m'}^{J_0^\kappa, I_0^\kappa} l_{m'}^{I^\kappa}$ for all box I^κ at level κ and all $1 \leq l' \leq d^\kappa$
 7. Compute $f_{\mathcal{F}, \text{FMM}}(\bar{x}_i) = \sum_{l'=1}^{d^\kappa} K(\bar{x}_i - c_{I^\kappa}, y_{l'}^{I_0^\kappa}) \hat{l}_{l'}^{I^\kappa}$, for all box I^κ and all $x_i \in I^\kappa$ at level κ
-

The complexity of the multilevel procedure is given in Table 2. For step 4 and the computation of the near interactions to be in the same order in N , one must impose $(2^D)^\kappa \propto N^2(2^D)^{-\kappa}$, leading to a overall complexity of the sum of order N . The choice that minimizes the overall cost is $\kappa \propto \log(N)$ and a constant number of points per box at the level κ . As for any multilevel FMM algorithm, the complexity of step 3 in Algorithm 2 has been reduced by replacing the sum over $J \in \mathcal{L}(I)$ (which is of complexity linear in N) by a sum over $J \in \mathcal{I}(I)$ (which is of complexity independent in N), making use of the fact that the set of boxes well-separated from I^k at level k contains the set of boxes well-separated from the parent of I^k at level $k - 1$. By precomputing some terms, independent to the location of the particles, the complexity of the steps 2, 4 and 5 can be reduced.

4 Optimizing the overall complexity

4.1 Precomputations for the M2M and L2L steps

Recall that when using an EIM approximation, there holds $\Delta = B^{-t}\Gamma^{-1}$, where the matrices B and Γ are triangular and constructed in a precomputation step. In practice, the condition number of the

Step number	Complexity
1	dN
2	$d^3(2^D)^\kappa$
3	$d^2(2^D)^\kappa$
4	$d^2 n_I (2^D)^\kappa$
5	$d^3(2^D)^\kappa$
6	$d^2(2^D)^\kappa$
7	dN
f_N	$N^2(2^D)^{-\kappa}$

Table 2: Complexity of each step of Algorithm 3 (for the most expensive level) and of the near interactions term ; n_I denotes the maximum number of elements in the interaction list of a box (4 in 1D, 40 in 2D and 316 in 3D).

matrix Γ worsens as the EIM approximation gets more accurate. It is well known in linear algebra that when a matrix Γ is ill-conditioned, a direct resolution of the linear system $\Gamma x = b$ for some $b \in \mathbb{R}^d$ provides a result much less polluted by numerical errors than first inverting the matrix Γ and then computing the matrix vector product $\Gamma^{-1}b$. Since the matrices Γ and B are triangular, computing $\Delta b = y$ by first solving $\Gamma x = b$ and then solving $B^t y = x$ is of the same complexity d^2 as the direct matrix-vector product.

Consider step 2 in Algorithm 3 at level k :

$$W_m^{I^k} = \sum_{p,q=1}^{d^{k+1}} \Delta_{p,q}^{I_0^{k+1}, J_0^{k+1}} \sum_{J^{k+1} \in \mathcal{C}(I^k)} K(x_m^{I_0^k} + c_{I^k} - c_{J^{k+1}}, y_p^{J_0^{k+1}}) W_q^{J^{k+1}}.$$

We see that $W_m^{I^k} = \sum_{J^{k+1} \in \mathcal{C}(I^k)} K^{I^k, J^{k+1}} \Delta_{m,p}^{I_0^{k+1}, J_0^{k+1}} W_q^{J^{k+1}}$, where $K_{m,p}^{I^k, J^{k+1}} = K(x_m^{I_0^k} + c_{I^k} - c_{J^{k+1}}, y_p^{J_0^{k+1}})$, $1 \leq m \leq d^k$, $1 \leq p \leq d^{k+1}$. The matrix $K_{M2M}^{I^k, J^{k+1}} := K^{I^k, J^{k+1}} \Delta_{0,0}^{I_0^{k+1}, J_0^{k+1}}$ can be precomputed as $((\Delta_0^{I_0^{k+1}, J_0^{k+1}})^t (K^{I^k, J^{k+1}})^t)^t$, where each matrix-vector product for each column of $(K^{I^k, J^{k+1}})^t$ is computed solving linear systems involving B and Γ matrices to preserve numerical accuracy, as explained in the previous paragraph. Notice that $K_{M2M}^{I^k, J^{k+1}}$ only depends on the relative position of I^k and J^{k+1} . Since J^{k+1} is a child of I^k , there are actually only 2^D different operators $K_{M2M}^{I^k, J^{k+1}}$, where D denotes the dimension of the space containing the particles. We denote them K_{M2M}^i , $1 \leq i \leq 2^D$, and step 2 in Algorithm 3 at level k is reduced to

$$W_m^{I^k} = \sum_{i=1}^{2^D} K_{M2M}^i W_q^{J_i^{k+1}}, \quad (28)$$

where J_i^{k+1} denotes the i^{th} child of I^k . This step is now of complexity $d^2 2^D (2^D)^k$ at level k , and the leading order of step 2 in Algorithm 3 is $d^2 2^D (2^D)^{\kappa-1} = d^2 (2^D)^\kappa$.

The same optimization is done for step 5 in Algorithm 3, reducing the complexity of this step to the order $d^2 (2^D)^\kappa$.

4.2 Optimization of the M2L step

It is well known that the M2L step is the most expensive one. Numerous efforts have been made to reduce the execution time of this step without degrading the accuracy of the overall result, see in particular [25]. We use the optimization denoted SARcmp in [25]. It consists of a bivariate ACA applied to the collection of all the M2L operators at each level, followed by a truncated SVD. The bivariate ACA and truncated SVD is then applied for each M2L operator at each level to increase further the compression. The procedure is detailed for symmetric kernels for completeness of the presentation.

Consider a level k and the collection of all the M2L operators at this level. Consider a level k , a box I^k of the tree at this level k , and a box J^k in the interaction list of I^k . The vector $\delta := c_{J^k} - c_{I^k}$ only depends on the relative position of the boxes I^k and J^k . Therefore, there are only $n_I = 7^D - 3^D$ different M2L matrices $K_{i,j}^\delta := K(x_i^{J^k}, y_j^{J^k} + \delta)$ per level (we recall that D denotes the dimension of the space containing the particles), see (19). We suppose that EIM algorithms have been carried-out for a certain error bound ϵ , leading to the choice of d_k interpolation points. The M2L matrices are organized in one large matrix as follows $K_{\text{fat}} := [K^{\delta_1} K^{\delta_2} \dots K^{\delta_{n_I}}]$, where $K_{\text{fat}} \in \mathbb{R}^{d_k \times n_I d_k}$. Two rectangular matrices U, V are computed by a bivariate ACA such that $|UV|_F \leq \epsilon |K_{\text{fat}}|_F$, where $|\cdot|_F$ denotes the Frobenius norm. Then, two QR decompositions are computed, such that $U = Q_U R_U$ and $(V)^t = Q_V R_V$, and a SVD is computed on $R_U (R_V)^t$. Denote s_i , $1 \leq i \leq d_k$, the singular values of $R_U (R_V)^t$ in decreasing order, and \hat{U}, \hat{V} the unitary matrices such that $R_U (R_V)^t = \hat{U} \text{diag}(s) \hat{V}$. The two following filters are applied: (i) keep the first j singular values such that $\frac{s_j}{s_0} \leq \epsilon \leq \frac{s_{j-1}}{s_0}$, (ii) keep the first j singular values such that $\frac{\sum_{i=1}^j s_i}{\sum_{i=1}^{d_k} s_i} \geq 1 - \epsilon \geq \frac{\sum_{i=1}^{j-1} s_i}{\sum_{i=1}^{d_k} s_i}$. We denote r_k the number of kept singular values. Denote \hat{U}_{r_k} and \hat{V}_{r_k} the first r_k columns of respectively \hat{U} and \hat{V} . Following [13], we can show that the M2L step can be reduced to size $r_k \times r_k$ (instead of size $d_k \times d_k$) matrix-vector products involving the matrices $C^\delta := \hat{U}_{r_k}^T K^\delta \hat{U}_{r_k}$, plus low complexity pre- and post-processing steps.

Then, the same compression is applied locally for each compressed M2L operator C^δ . This time, the bivariate ACA is applied on each C^δ , producing the rectangular matrices U and V . The QR decompositions, SVD and filters are applied in the same fashion, selecting the singular values s_k . The rectangular matrices \hat{U}_{s_k} and \hat{V}_{s_k} are constructed, and a diagonal matrix S_{s_k} , whose diagonal entries are the square roots of the selected singular values, is constructed. Then, we construct $\tilde{U}_{s_k} = Q_U \hat{U}_{s_k} S_{s_k} \in \mathbb{R}^{r_k \times s_k}$ and $\tilde{V}_{s_k} = S \hat{V}_{s_k} Q_V^t \in \mathbb{R}^{s_k \times r_k}$, so that C^δ is approximated by $\tilde{U}_{s_k} \tilde{V}_{s_k}$. Thus, size $r_k \times r_k$ matrix-vector products involving the matrices C^δ are replaced by two $r_k \times s_k$ matrix-vector products involving the matrices \tilde{U}_{s_k} and \tilde{V}_{s_k} .

5 Numerical experiments in 3D

5.1 Some elements on the implementation

The EIFMM has been implemented in C++, using the open source library `scalfm`, see [9]. From [29], `scalfm` is one of the fastest opensource library for FMM computations. We will mainly compare our implementation of the EIFMM to the Chebyshev interpolation-based implementation in `scalfm`, which follows the Black-Box FMM [13]. In our code, the EIM precomputation and the SARcmp precomputation steps are performed in C++ using BLAS and LAPACK, and stored in binary files. These precomputations are then read by the modified `scalfm` library, using also BLAS for the basic linear algebra operations.

5.2 Test-cases

We consider sets of points included in $\Omega = (-0.5, 0.5)^3$, randomly taken following (i) a uniform law in Ω , (ii) a uniform law on the sphere of radius 0.5 and (iii) a modification of the set (ii) in an ellipsoid, see Figure 5.

When comparing runtimes, we have to make sure that the Black-Box FMM is used in a optimal way. There are two parameters to choose: the number of interpolation points per dimension and the accuracy ϵ_c of the compression. First, we chose 3 interpolation points per dimension and compute the relative error with $\epsilon_c = 10^{-15}$. Then, if this relative error is smaller than the desired one, we search by dichotomy the largest value of the ϵ_c so that the relative error is smaller or equal to the desired one. Otherwise, we increase the number of interpolation points per dimension by one and repeat the process. That way, the runtimes comparison is fair in the sense that Block-Box FMM is not artificially slowed down by a too large number of interpolation points together with a too large value for ϵ_c . Notice that this trick is necessary since we do not know in advance how the Chebyshev interpolation will perform in practice for a new kernel, while for the EIFMM, the error estimate automatically selects the right number of interpolation points.




Figure 5: Sets of points considered in the numerical applications, from left to right: random (5000 points), sphere (10100 points), and ellipsoid (1944 points). The values for the number of points have been chosen for representation reasons, higher values will be considered in the numerical experiments.

5.3 Performance for the Laplace kernel

In this section, we compare runtimes of the FMM summation using the Black-Box FMM and the EIFMM. The computations are carried out in sequential, on a laptop with 4 GB of RAM and a Intel(R) Core(TM) i5 CPU M 560 @ 2.67GHz. For the Black-Box FMM, we use the so-called symmetric M2L compression, based on local SVD for each M2L operator. The symmetric M2L compression exploits the symmetry of the Chebyshev nodes to express the M2L step as matrix-matrix products, and even if the complexity of this step is not improved, runtimes are drastically improved due to better reuse of memory cache by the processor. Proposed in [25], the symmetric compression is currently, to the author's knowledge, the most efficient compression strategy for the Black-Box FMM. Since the approximation (9) used in EIFMM has four summations instead of two for Black-Box FMM, the steps 3 and 6 in Algorithm 3 for EIFMM are not present in Black-Box FMM.

In Figure 6 are represented execution time with respect to the relative error on the FMM summation for the Laplace kernel, for the symmetric Black-Box FMM, the EIFMM and the fastest algorithm available in scalfm. This latter is dedicated to the Laplace kernel, explaining why it performs better than the other two. The test-case corresponds to 10^6 random points taken in the unit cube, organized in a 6-level tree, see Figure 5, left picture. The execution times are better for Black-Box FMM at low accuracy, and equivalent between EIFMM and Black-Box FMM at high accuracy. Notice that since the Laplace kernel is homogeneous, the difficulty of the approximation problem does not change from one level to another in the tree. As a consequence, the variable-order capability of our algorithm is not exploited for this kernel. In what follows, inhomogeneous kernels are considered and gains on runtimes are observed. Precomputation times will be investigated in the following section as well.

5.4 Performance for some inhomogeneous kernels

The advantages of the EIFMM mainly apply to inhomogeneous kernels, for which the error estimation allows the selection of an optimal number of interpolation points at each level of the tree by the EIM.

5.4.1 1-million-point test-cases

The three test-cases presented in Figure 5 are considered: 10^6 random points in the unit cube with a 6-level tree (denoted "cube"), 992,046 points at the surface of a sphere and an ellipsoid, with respectively 8-level and 9-level trees (denoted respectively "sphere" and "ellipsoid").

Figures 7, 8 and 9 present execution times of the long-range FMM accelerated terms with respect to the relative error on the FMM summation, respectively for the kernels $\frac{\cos 20r}{r}$, e^{-r^2} and $\sqrt{r^2 + 1}$, for the three test-cases (cube, sphere and ellipsoid). The reported runtimes do not contain the times required to build the tree, nor the EIM, M2L compression, M2M and L2L precomputations. For the EIFMM, we use the SARcmp compression for the M2L operators, as described in Section 4.2. Due to smaller interpolation formulae, EIFMM enables memory savings with respect to Black-Box FMM. In the figures, missing data correspond to cases requiring more memory than available with our computer,




Figure 6: Execution time with respect to the relative error on the FMM summation for the Laplace kernel; black: symmetric Black-Box FMM, blue: EIFMM, red: fastest spherical algorithm available in scalfmm.

for which EIFMM can be used, but not Black-Box FMM. For the three kernels, and especially the last two, the EIFMM is much faster than the Black-Box FMM, because the error estimation and the greedy procedure in EIFMM allows the selection of an optimal number of interpolation points at each level of the tree. For instance, for the kernel e^{-r^2} on the ellipsoid and a relative error of 10^{-6} , the Black-Box FMM needs 6 interpolation points per dimension (i.e. 216 interpolation points), whereas the EIFMM needs respectively 167, 74, 36, 21, 13, 10, 9, and 4 points at levels 2 to 9 of the tree. As we go deeper in the tree, the interpolation problems become easier, and the EIM selects fewer points to produce an interpolation of the same quality.

Table 3 provides precomputations and executions times for the kernels $\frac{\cos(20r)}{r}$, e^{-r^2} and $\sqrt{r^2 + 1}$ on the cube test-case. Notice that in order not to artificially disadvantage Black-Box FMM, we report values with EIFMM set to a better relative accuracy and we neglect the precomputation times of Black-Box FMM (which are short anyway). In the reported cases, EIFMM becomes advantageous (i.e. the EIFMM precomputations times are compensated by faster long-range FMM terms computation time) if one has to compute more than 2 to 18 FMM summations.

kernel	$\frac{\cos(20r)}{r}$		e^{-r^2}		$\sqrt{r^2 + 1}$	
algorithm	EIFMM	BBFMM	EIFMM	BBFMM	EIFMM	BBFMM
relative accuracy	2×10^{-4}	4×10^{-4}	2×10^{-6}	3×10^{-6}	1×10^{-5}	2×10^{-5}
EIM precomp time (s)	73	—	30	—	10	—
compression time (s)	177	—	33	—	1.1	—
tree construction time (s)	3		3		3	
long-range FMM terms comp time (s)	52	69	5.4	37	3.5	4.3
short-range terms comp time(s)	32		30		16	
minimum EIFMM calls	18		2		14	

Table 3: Precomputations and executions times for the kernels $\frac{\cos(20r)}{r}$, e^{-r^2} and $\sqrt{r^2 + 1}$ on the cube test-case. The last line indicates the minimum FMM summations to compute for EIFMM to be faster than Black-Box FMM (BBFMM), taking all precomputations times into account.

5.4.2 100-million-point test-case

We consider 10^8 random points in the unit cube (Figure 5, left) in a 9-level tree, with the e^{-r^2} kernel. The computations are carried-out in sequential on a Intel(R) Xeon(R) CPU E5-4620 v2 @ 2.60GHz

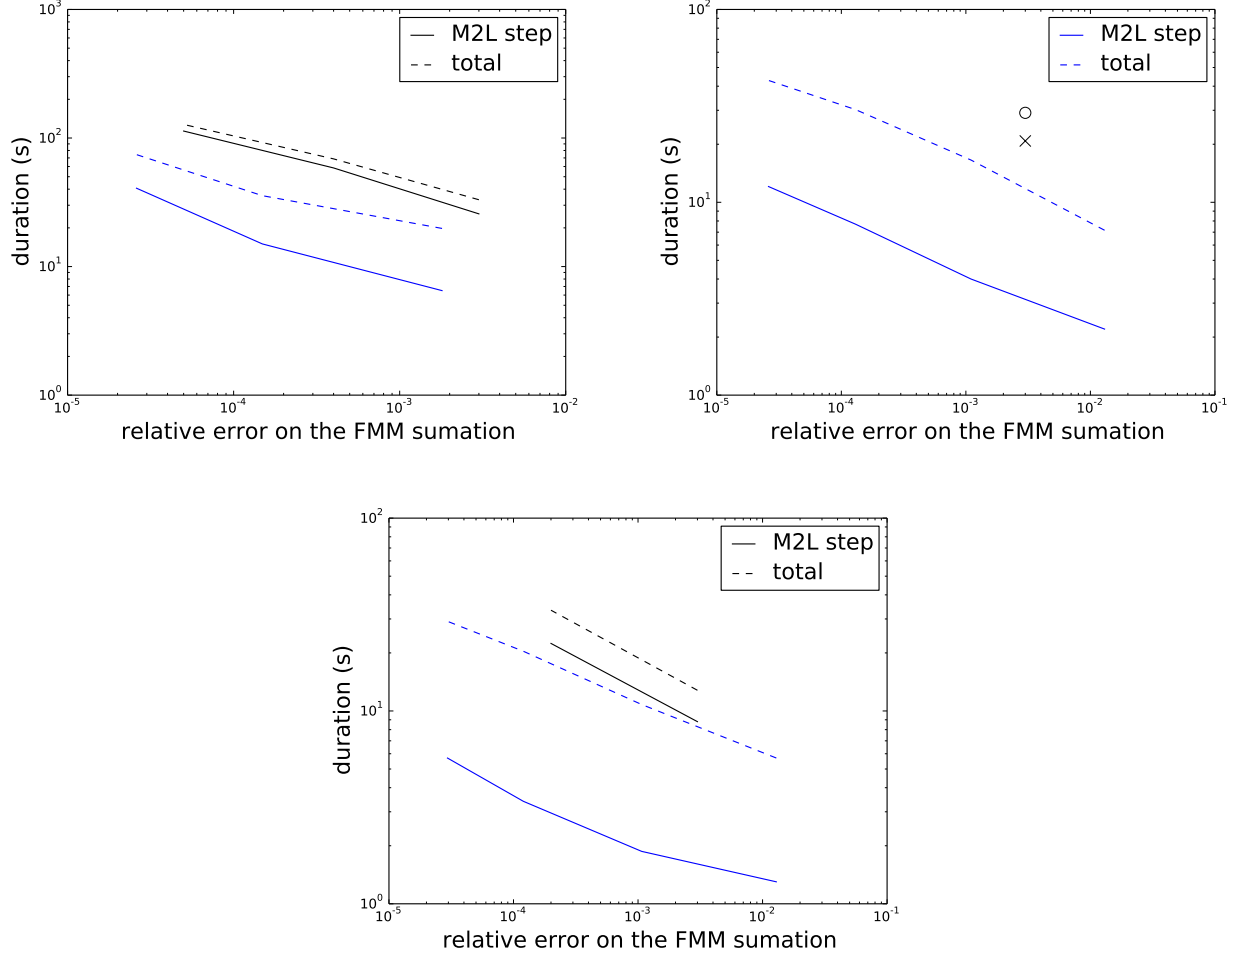


Figure 7: Execution time of the long-range FMM accelerated terms with respect to the relative error on the FMM summation for the kernel $\frac{\cos 20r}{r}$, black: symmetric Black-Box FMM, blue: EIFMM. From top to bottom and left to right: cube, sphere and ellipsoid test-cases.

and 128 GB of RAM.

Table 4 presents the relative accuracy and the execution times of all the steps for two simulations using EIFMM and Black-Box FMM. For a better accuracy, the EIFMM is faster than the Black-Box FMM, including the precomputation times (increasing the accuracy with Black-Box FMM would require more than 128 GB of RAM). In this case, EIFMM is interesting even for a single FMM computation. Notice that the precomputations for EIFMM are negligible for such large cases: the precomputation complexity is linear with respect to the number of levels in the tree (therefore logarithmic with respect to N). Moreover, for inhomogeneous kernels, the approximation problem for an additional level is easier than the previous approximation problems, leading to even faster additional EIM and M2L compression precomputation times.

6 Conclusion and outlook

This work introduces a new Fast Multipole Method (FMM) using a low-rank approximation of the kernel based on the Empirical Interpolation Method (EIM), called the Empirical Interpolation Fast Multipole Method (EIFMM). The proposed multilevel algorithm is implemented in `scalnm`, a FMM library written in C++. The important feature of the EIFMM is a built-in error estimation of the interpolation error made by the low-rank approximation of the far-field behavior of the kernel. As a consequence, the algorithm automatically selects the optimal number of interpolation points required to ensure a given accuracy for the result, leading to important execution time reduction for

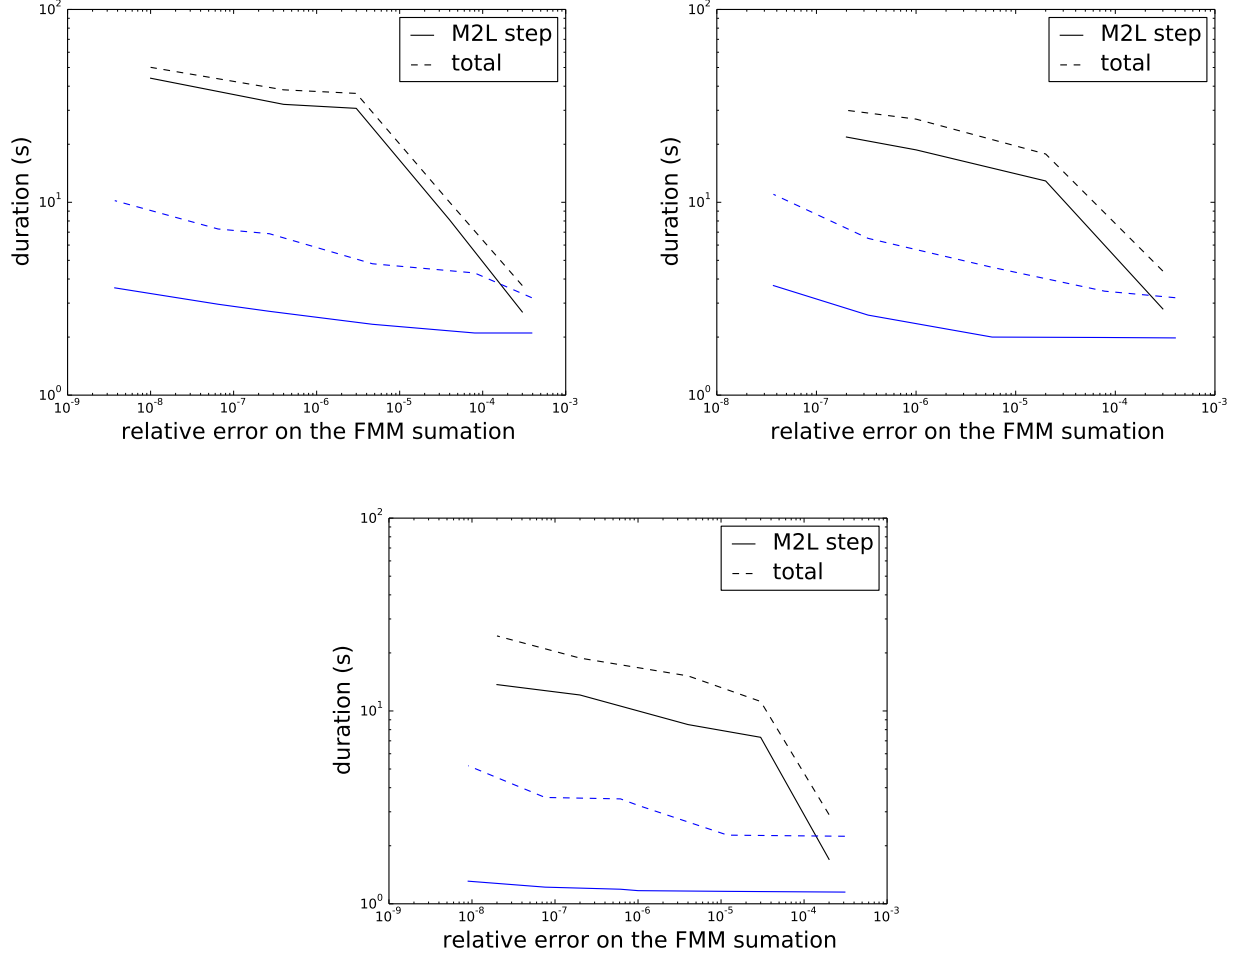


Figure 8: Execution time of the long-range FMM accelerated terms with respect to the relative error on the FMM summation for the kernel e^{-r^2} , black: symmetric Black-Box FMM, blue: EIFMM. From top to bottom and left to right: cube, sphere and ellipsoid test-cases.

inhomogeneous kernels, for which the difficulty of the approximation varies from one level of the tree to another.

Acknowledgement

The author wish to thank Centre National d'Etudes Spatiales for financial support through the TOSCA committee, Guillaume Sylvand (Airbus Group Innovations, INRIA) for drawing his attention to fast multipole methods, thoroughly reading the manuscript and for many fruitful discussions, Tony Lelièvre (CERMICS) and Alexandre Ern (CERMICS) for fruitful discussions, and Mathieu Ruffat (Diginext) for his C++ expertise.

References

- [1] B. Alpert, G. Beylkin, R. Coifman, and V. Rokhlin. Wavelet-like bases for the fast solutions of second-kind integral equations. *SIAM J. Sci. Comput.*, 14(1):159–184, 1993.
- [2] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathématique*, 339(9):667 – 672, 2004.

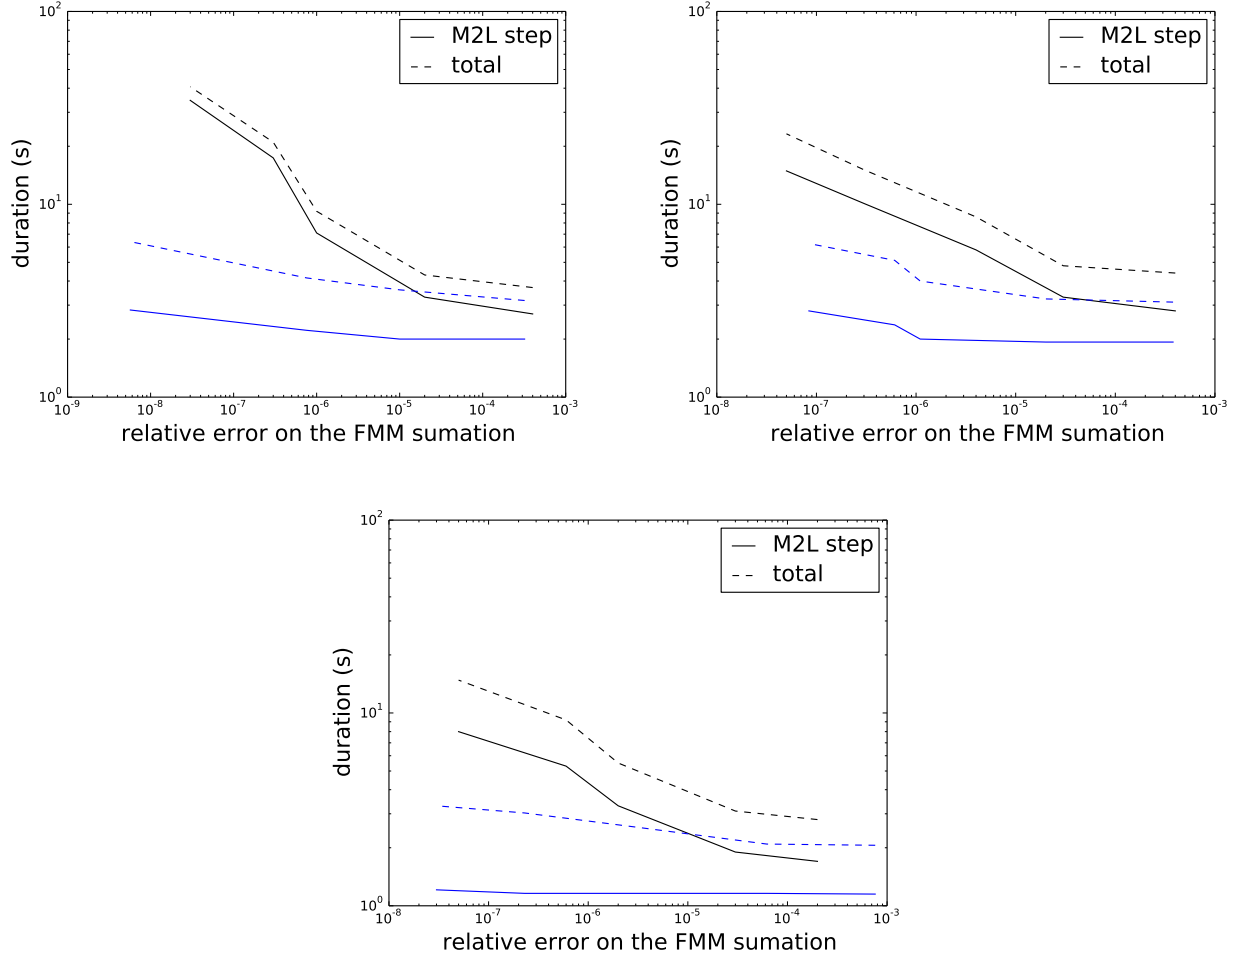


Figure 9: Execution time of the long-range FMM accelerated terms with respect to the relative error on the FMM summation for the kernel $\sqrt{r^2 + 1}$, black: symmetric Black-Box FMM, blue: EIFMM. From top to bottom and left to right: cube, sphere and ellipsoid test-cases.

- [3] M. Bebendorf. Approximation of boundary element matrices. *Numerische Mathematik*, 86(4):565–589, 2000.
- [4] M. Bebendorf and S. Rjasanow. Adaptive low-rank approximation of collocation matrices. *Computing*, 70(1):1–24, 2003.
- [5] M. Bebendorf and R. Venn. Constructing nested bases approximations from the entries of non-local operators. *Numerische Mathematik*, 121(4):609–635, 2012.
- [6] S. Börm, M. Löhndorf, and J.M. Melenk. Approximation of integral operators by variable-order interpolation. *Numerische Mathematik*, 99(4):605–643, 2005.
- [7] F. Casenave, A. Ern, and T. Lelièvre. A nonintrusive reduced basis method applied to aeroacoustic simulations. *Advances in Computational Mathematics*, pages 1–26, 2014.
- [8] H. Cheng, Z. Gimbutas, P.G. Martinsson, and V. Rokhlin. On the compression of low-rank matrices. *SIAM J. Sci. Comput.*, 26(4):1389–1404, 2005.
- [9] O. Coulaud, B. Bramas, and C. Piacibello. Scalffmm, C++ Fast Multipole Method Library for HPC. <http://scalffmm-public.gforge.inria.fr/doc/>.
- [10] A. Dutt, M. Gu, and V. Rokhlin. Fast algorithms for polynomial interpolation, integration, and differentiation. *SIAM Journal on Numerical Analysis*, 33(5):1689–1711, 1996.

algorithm	EIFMM	Black-Box FMM
relative accuracy	7×10^{-5}	6×10^{-4}
EIM precomp time	22.8 s	–
compression time	1.3 s	–
tree construction time (s)	9 min 10 s	
P2M (1)	1 min 13 s	2 min 17 s
M2M (2-3)	31 s	49 s
M2L (4)	10 min 39 s	53 min 01 s
L2L (5-6)	41 s	59 s
L2P (7)	1 min 42 s	3 min 55 s
total long-range FMM terms	14 min 46 s	1 h 01 min 01 s
short-range terms	19 min 36 s	19 min 36 s
total	43 min 56 s	1 h 29 min 47 s

Table 4: Accuracy and execution times for the 100 million points test-case. P2M, M2M, M2L, L2L and L2P corresponds respectively to the steps (1), (2-3), (4), (5-6) and (7) in Algorithm 3.

- [11] A. Dutt and V. Rokhlin. Fast fourier transforms for nonequispaced data. *SIAM Journal on Scientific Computing*, 14(6):1368–1393, 1993.
- [12] J. L. Eftang and B. Stamm. Parameter multi-domain ‘hp’ empirical interpolation. *International Journal for Numerical Methods in Engineering*, 90:412–428, 2012.
- [13] W. Fong and E. Darve. The black-box fast multipole method. *Journal of Computational Physics*, 228(23):8712 – 8725, 2009.
- [14] K. Giebermann. Multilevel approximation of boundary integral operators. *Computing*, 67(3):183–207, 2001.
- [15] Z. Gimbutas, L. Greengard, and M. Minion. Coulomb interactions on planar structures: Inverting the square root of the laplacian. *SIAM Journal on Scientific Computing*, 22(6):2093–2108, 2001.
- [16] S.A. Goreinov, E.E. Tyrtyshnikov, and N.L. Zamarashkin. A theory of pseudoskeleton approximations. *Linear Algebra and its Applications*, 261(1–3):1 – 21, 1997.
- [17] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *Journal of Computational Physics*, 135(2):280–292, 1997.
- [18] W. Hackbusch. A sparse matrix arithmetic based on H-matrices. part I: Introduction to H-matrices. *Computing*, 62(2):89–108, 1999.
- [19] W. Hackbusch and S. Börm. Data-sparse approximation by adaptive H2-matrices. *Computing*, 69(1):1–35, 2002.
- [20] W. Hackbusch and S. Börm. H2-matrix approximation of integral operators by interpolation. *Applied Numerical Mathematics*, 43(1–2):129 – 143, 2002. 19th Dundee Biennial Conference on Numerical Analysis.
- [21] W. Hackbusch and Z.P. Nowak. On the fast matrix multiplication in the boundary element method by panel clustering. *Numerische Mathematik*, 54(4):463–491, 1989.
- [22] Y. Maday and O. Mula. A generalized empirical interpolation method: Application of reduced basis techniques to data assimilation. 4:221–235, 2013.
- [23] Y. Maday, N. C. Nguyen, A. T. Patera, and S. Pau. A general multipurpose interpolation procedure: the magic points. *Communications On Pure And Applied Analysis*, 8(1):383–404, 2008.

- [24] P.G. Martinsson and V. Rokhlin. An accelerated kernel-independent fast multipole method in one dimension. *SIAM J. Sci. Comput.*, 26:1389–1404, 2005.
- [25] M. Messner, B. Bramas, D. Coulaud, and E. Darve. Optimized M2L Kernels for the Chebyshev Interpolation based Fast Multipole Method. *CoRR*, abs/1210.7292, 2012.
- [26] B. Peherstorfer, D. Butnaru, K. Willcox, and H.-J. Bungartz. Localized discrete empirical interpolation method. *SIAM Journal on Scientific Computing*, 36(1):A168–A192, 2014.
- [27] E. Tyrtysnikov. Mosaic-skeleton approximations. *CALCOLO*, 33(1-2):47–57, 1996.
- [28] L. Ying, G. Biros, and D. Zorin. A kernel-independent adaptive fast multipole algorithm in two and three dimensions. *J. Comput. Phys.*, 196(2):591–626, 2004.
- [29] R. Yokota. Fast multipole methods benchmark. <https://sites.google.com/site/rioyokota/research/fmm>.